arXiv:1608.03928v1 [math.OC] 13 Aug 2016

Hybrid Jacobian and Gauss-Seidel proximal block coordinate update methods for linearly constrained convex programming

Yangyang Xu*

Abstract

Recent years have witnessed the rapid development of block coordinate update (BCU) methods, which are particularly suitable for problems involving large-sized data and/or variables. In optimization, BCU first appears as the coordinate descent method that works well for smooth problems or those with separable nonsmooth terms and/or separable constraints. As nonseparable constraints exist, BCU can be applied under primal-dual settings.

In the literature, it has been shown that for weakly convex problems with nonseparable linear constraint, BCU with fully Gauss-Seidel updating rule may fail to converge and that with fully Jacobian rule can converge sublinearly. However, empirically the method with Jacobian update is usually slower than that with Gauss-Seidel rule. To maintain their advantages, we propose a hybrid Jacobian and Gauss-Seidel BCU method for solving linearly constrained multi-block structured convex programming, where the objective may have a nonseparable quadratic term and separable nonsmooth terms. At each primal block variable update, the method approximates the augmented Lagrangian function at an affine combination of the previous two iterates, and the affinely mixing matrix with desired nice properties can be chosen through solving a semidefinite programming. We show that the hybrid method enjoys the theoretical convergence guarantee as Jacobian BCU. In addition, we numerically demonstrate that the method can perform as well as Gauss-Seidel method and better than a recently proposed randomized primal-dual BCU method.

Keywords: block coordinate update (BCU), Jacobian rule, Gauss-Seidel rule, alternating direction method of multipliers (ADMM)

Mathematics Subject Classification: 9008, 90C25, 90C06, 68W40.

1 Introduction

Driven by modern applications in image processing, statistical and machine learning, block coordinate update (BCU) methods have revived in recent years. BCU methods decompose a complicated

*yangyang.xu@ua.edu. Department of Mathematics, University of Alabama, Tuscaloosa, AL 35487

large-scale problem into easy small subproblems and tend to have low per-update complexity and low memory requirement, and they give rise to powerful ways to handle problems involving large-sized data and/or variables. These methods originate from the coordinate descent method that only applies to optimization problems with separable constraints. Under primal-dual settings, they have been developed to deal with nonseparably constrained problems.

In this paper, we consider the linearly constrained multi-block structured problem with a quadratic term in the objective:

$$\min_x \frac{1}{2}x^\top Qx + d^\top x + \sum_{i=1}^m g_i(x_i),\ \text{s.t. } Ax = b, \tag{1}$$

where the variable is partitioned into m disjoint blocks $x = (x_1, x_2, \ldots, x_m)$, Q is a positive semidefinite (PSD) matrix, and each g_i is a proper closed convex and possibly non-differentiable function. Note that part of g_i can be an indicator function of a convex set $\mathcal{X}_i$, and thus (1) can implicitly include certain separable block constraints $x_i \in \mathcal{X}_i$ in addition to the nonseparable linear constraint.

Corresponding to the partition of x, we let $d = (d_1, \ldots, d_m)$ and define $\tilde{g}_i(x_i) = g_i(x_i) + d_i^\top x_i$, $\forall i$. This way, $d^\top x$ is absorbed into g_i's. Hence, without loss of generality, we assume $d = 0$. Due to its multi-block and also coordinate friendly [33] structure, we will derive a BCU method for solving (1), by performing BCU to x_i's based on the augmented Lagrangian function of (1), followed by an update to the Lagrangian multiplier; see the updates in (7) and (8).

1.1 Motivating examples

Many applications can be formulated in the form of (1). If g_i is the indicator function of the nonnegative orthant for all i, then it reduces to the nonnegative linearly constrained quadratic programming (NLCQP):

$$\min_x \frac{1}{2}x^\top Qx + d^\top x,\ \text{s.t. } Ax = b,\ x \geq 0. \tag{2}$$

All convex QPs can be written as NLCQP by adding slack variable and/or decomposing a free variable into positive and negative parts. As Q is a huge-sized matrix, it would be beneficial to partition it into block matrices and correspondingly x and A into block variables and block matrices, and then apply BCU methods toward finding a solution to (2).

Another example is the constrained Lasso regression problem proposed in [26]:

$$\min_x \frac{1}{2}\|Ax - b\|^2 + \mu\|x\|_1,\ \text{s.t. } Cx \leq d. \tag{3}$$

If there is no constraint $Cx \leq d$, (3) simply reduces to the Lasso regression problem [38]. Introducing a nonnegative slack variable y, we can write (3) into the form of (1):

$$\min_{x,y} \frac{1}{2}\|Ax - b\|^2 + \mu\|x\|_1 + \iota_+(y),\ \text{s.t. } Cx + y = d, \tag{4}$$

where $\iota_+(y)$ is the indicator function of the nonnegative orthant, equaling *zero* if y is nonnegative and $+\infty$ otherwise. Again for large-sized A or C, it is preferable to partition x into disjoint blocks and apply BCU methods to (4).

There are many other examples arising in signal and image processing and machine learning such as the compressive principal component pursuit [45] (see (49) below) and the regularized multiclass support vector machines [48] (see (51) below for instance). More examples, including basis pursuit, conic programming, and the exchange problem, can be found in [11,19,23,28,37] and the references therein.

1.2 Related works

BCU methods in optimization first appear in [22] as the coordinate descent (CD) method for solving quadratic programming with separable nonnegative constraints but without nonseparable equality constraint. The CD method updates one coordinate every time while all the remaining ones are fixed. It may stuck at a non-stationary point if there are nonseparable nonsmooth terms in the objective; see the example in [44]. On solving smooth problems or those with separable nonsmooth terms, the convergence properties of the CD method have been intensively studied (e.g., [24, 36, 40, 42, 46, 50, 51]). For the linearly constrained problem (1), the CD method can also stuck at a non-stationary point, for example, if the linear constraint is simply $x_1 - x_2 = 0$. Hence, to directly apply BCU methods to linearly constrained problems, at least two coordinates need be updated every time; see [32, 41] for example.

Another way of applying BCU towards finding a solution to (1) is to perform primal-dual block coordinate update (e.g., [7,33–35]). These methods usually first formulate the first-order optimality system of the original problem and then apply certain operator splitting methods. Assuming monotonicity of the iteratively performed operator (that corresponds to convexity of the objective), almost sure iterate sequence convergence to a solution can be shown, and with strong monotonicity assumption (that corresponds to strong convexity), linear convergence can be established.

In the literature, there are also plenty of works applying BCU to the augmented Lagrangian function as we did in (7). One popular topic is the alternating direction method of multipliers (ADMM) applied to separable multi-block structured problems, i.e., in the form of (1) without the nonseparable quadratic term. Originally, ADMM was proposed for solving separable two-block problems [13, 16] by cyclicly updating the two block variables in a Gauss-Seidel way, followed by an update to the multiplier, and its convergence and also $O(1/t)$ sublinear rate is guaranteed by assuming merely weak convexity (e.g., [21, 30, 31]). While directly extended to problems with more than two blocks, ADMM may fail to converge as shown in [4] unless additional assumptions are made such as strong convexity on part of the objective (e.g., [2, 6, 10, 17, 27, 29, 30]), orthogonality condition on block matrices in the linear constraint (e.g., [4]), and Lipschitz differentiability of the objective and invertibility of the block matrix in the constraint about the last updated block variable (e.g., [43]). For problems that do not satisfy these conditions, ADMM can be modified and have guaranteed

convergence and even rate estimate by adding extra correction steps (e.g., [19,20]) or by randomized block coordinate update (e.g., [9,14]) or adopting Jacobian updating rules (e.g., [11,18,28,37]) that essentially reduce the method to proximal augmented Lagrangian or two-block ADMM method.

When there is a nonseparable term coupling variables together in the objective like (1), existing works usually replace the nonseparable term by a relatively easier majorization function during the iterations and perform the upper-bound or majorized ADMM updates. For example, [23] considers generally multi-block problems with nonseparable Lipschitz-differentiable term. Under certain error bound conditions and a diminishing dual stepsize assumption, it shows subsequence convergence. Along a similar direction, [8] specializes the method of [23] to two-block problems and establishes global convergence and also $O(1/t)$ rate with any positive dual stepsize. Without changing the nonseparable term, [15] adds proximal terms into the augmented Lagrangian function during each update and shows $O(1/t)$ sublinear convergence by assuming strong convexity on the objective. Very recently, [5] directly applies ADMM to (1) with $m = 2$, i.e., only two block variables, and establishes iterate sequence convergence to a solution while no rate estimate has been shown.

1.3 Contributions

We summarize our contributions as follows.

- We propose a hybrid Jacobian and Gauss-Seidel BCU method for solving (1). Through affinely combining two most recent iterates, the proposed method can update block variables, in Jacobian or Gauss-Seidel manner, or a mixture of them, Jacobian rules within groups of block variables and a Gauss-Seidel way between groups. It can enjoy the theoretical convergence guarantee as BCU with fully Jacobian update rule and also practical fast convergence as the method with fully Gauss-Seidel rule.

- We establish global iterate sequence convergence and $O(1/t)$ rate results of the proposed BCU method, by assuming certain conditions on the affinely mixing matrix (see Assumption 2) and choosing appropriate weight matrix P_i's in (7). Different from the after-correction step proposed in [19] to obtain convergence of multi-block ADMM with fully Gauss-Seidel rule, the affine combination of two iterates we propose can be regarded as a correction step before updating every block variable. In addition, our method allows nonseparable quadratic terms in the objective while the method in [19] can only deal with separable terms. Furthermore, utilizing prox-linearization technique, our method can have much simpler updates.

- We discuss how to choose the affinely mixing matrix, which can be determined through solving a semidefinite programming (SDP); see (43). One can choose a desired matrix by adding certain constraints in the SDP. Compared to the original problem, the SDP is much smaller and can be solved offline. We demonstrate that the algorithm with the combining matrix found in this way can perform significantly better than that with all-one matrix (i.e., with fully Jacobian rule).

- We test the proposed BCU method on quadratic programming, the compressive principal component pursuit (see (48) below), and the multi-class support vector machine (see (51) below). By adapting P_i's in (7), we demonstrate that our method can outperform a randomized BCU method recently proposed in [14] and is comparable to the direct ADMM that has no guaranteed convergence.

1.4 Outline

The rest of the paper is organized as follows. In section 2, we present our algorithm, and we show its convergence with sublinear rate estimate in section 3. In section 4, we discuss how to choose the affinely mixing matrices used in the algorithm. Numerical experiments are performed in section 5. Finally, section 6 concludes the paper.

2 Algorithm

In this section, we present a BCU method for solving (1). First, we formulate (1) in an equivalent form for the convenience of our description and the convergence analysis later. Since Q is PSD, there always exists a matrix H such that $Q = H^\top H$. This factorization is not required in the implementation of our algorithm but only used in the convergence analysis. Corresponding to x, we partition A and H into m block matrices:

$$A = (A_1, \ldots, A_m), \quad H = (H_1, \ldots, H_m).$$

Letting

$$f(x) = \frac{1}{2} x^\top Q x = \frac{1}{2}\|Hx\|^2,$$

and recalling $d = 0$, we write (1) into

$$\min_x f(x) + \sum_{i=1}^m g_i(x_i), \text{ s.t. } \sum_{i=1}^m A_i x_i = b. \tag{5}$$

Now we derive our algorithm by applying a cyclic proximal block coordinate update to the augmented Lagrangian function of (5), that is

$$\mathcal{L}_\beta(x, \lambda) = f(x) + \sum_{i=1}^m g_i(x_i) - \langle \lambda, Ax - b\rangle + \frac{\beta}{2}\|Ax - b\|^2. \tag{6}$$

At each iteration, we first renew every x_i by minimizing a function, which approximates $\mathcal{L}_\beta$ at points that affinely combine the previous two iterates, one primal block variable at a time while all the remaining ones are fixed, and then update the multiplier λ by a dual gradient ascent step. More

precisely, at each iteration k, we perform the following cyclic BCU to x_i's and λ. For $i = 1, \ldots, m$, do the updates

$$x_i^{k+1} \in \arg\min_{x_i} \left\langle \nabla_i f(\hat{x}^{k,i}) - A_i^\top \left(\lambda^k - \beta(A\tilde{x}^{k,i} - b)\right), x_i \right\rangle + g_i(x_i) + \frac{1}{2}\|x_i - x_i^k\|_{P_i}^2, \tag{7}$$

and then

$$\lambda^{k+1} = \lambda^k - \rho(Ax^{k+1} - b), \tag{8}$$

where each P_i is a PSD matrix, and

$$\hat{x}_j^{k,i} = x_j^{k+1} - \hat{w}_{ij}(x_j^{k+1} - x_j^k), \quad \tilde{x}_j^{k,i} = x_j^{k+1} - \tilde{w}_{ij}(x_j^{k+1} - x_j^k), \ \forall i, j. \tag{9}$$

Here, $\hat{W}$ and $\tilde{W}$ are two (possibly different) mixing matrices. Practically, two common choices are

$$\hat{W} = \tilde{W} = \begin{bmatrix} 1 & 1 & \cdots & 1 \\ 1 & 1 & \cdots & 1 \\ \vdots & \ddots & \ddots & \vdots \\ 1 & \cdots & 1 & 1 \end{bmatrix}, \text{ and } \hat{W} = \tilde{W} = \begin{bmatrix} 1 & 1 & \cdots & 1 \\ 0 & 1 & \cdots & 1 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & 1 \end{bmatrix},$$

corresponding to fully Jacobian and fully Gauss-Seidel updates respectively, and the latter one usually performs better. However, theoretically, the latter case may fail to converge when $m \geq 3$, as shown in [4, 12]. Therefore, we will design mixing matrices between the above two choices to theoretically guarantee convergence and also main practically nice performance. Our choice is inspired from the convergence analysis and will be discussed in details in section 4.

Algorithm 1 summarizes the proposed method.

Algorithm 1: Hybrid Jacobian and Gauss-Seidel proximal block coordinate update for (5)

Initialization: choose x^1 and set $\lambda^1 = 0$; choose parameters $\beta, \rho, \hat{W}, \tilde{W}$ and P.
for $k = 1, 2, \ldots$ **do**
  for $i = 1, 2, \ldots, m$ **do**
    Update $x_i^{k+1} \leftarrow$ (7).
  Update $\lambda^{k+1} \leftarrow$ (8).
  if *a certain stopping condition is satisfied* **then**
    Return (x^{k+1}, λ^{k+1})

A few remarks on the algorithm are in place. First, we linearize the quadratic function f and the augmented term $\frac{\beta}{2}\|Ax - b\|^2$ at possibly different points. This way, we can explore different structures of Q and A. Secondly, to make x_i^{k+1} independent of later updated block variables, we will enforce the strictly upper triangular parts of the mixing matrices $\hat{W}$ and $\tilde{W}$ to all *one*'s. How to choose them with nice properties will be discussed in section 4. Thirdly, P_i can include $H_i^\top H_i$

(resp. $A_i^\top A_i$), which indicates if to linearize the quadratic term f (resp. the augmented term $\|Ax-b\|^2$) at x_i^k. If $P_i = \eta_i I + H_i^\top H_i$, then (7) reduces to

$$x_i^{k+1} \in \arg\min_{x_i} f(x_i, \hat{x}_{\neq i}^{k,i}) - \left\langle A_i^\top \left(\lambda^k - \beta(A\tilde{x}^{k,i} - b)\right), x_i \right\rangle + g_i(x_i) + \frac{\eta_i}{2}\|x_i - x_i^k\|^2, \tag{10}$$

and similarly one can get the update rule for $P_i = \eta_i I + \beta A_i^\top A_i$. Depending on the applications, one can choose appropriate P_i's to make the subproblem in (7) easy to solve and also maintain overall fast convergence; see the examples in section 5.

3 Convergence analysis

In this section, we analyze the convergence of Algorithm 1. We establish its global iterate sequence convergence and $O(1/t)$ rate by choosing appropriate mixing matrices $\hat{W}$ and $\tilde{W}$ and assuming merely weak convexity on the problem.

3.1 Notation and preliminary results

Before proceeding with our analysis, we introduce some notation and a few preliminary lemmas.

We let

$$g(x) = \sum_{i=1}^m g_i(x_i), \quad F = f + g.$$

A point x^* is a solution to (5) if there exists λ^* such that the KKT conditions hold:

$$0 \in \partial F(x^*) - A^\top \lambda^*, \tag{11a}$$

$$Ax^* - b = 0, \tag{11b}$$

where ∂F denotes the subdifferential of F. Together with the convexity of F, (11) implies

$$F(x) - F(x^*) - \langle \lambda^*, Ax - b\rangle \geq 0, \ \forall x. \tag{12}$$

We denote $\mathcal{X}^*$ as the solution set of (5). For any vector v and any symmetric matrix W of appropriate size, we define $\|v\|_W^2 = v^\top W v$. Note this definition does not require W to be PSD, so $\|v\|_W^2$ may be negative. I is reserved for the identity matrix and E for the all-one matrix, whose size would be clear from the context. $A \otimes B$ represents the Kronecker product of two matrices A and B. For any matrices A, B, C and D of appropriate sizes, it holds that (c.f., [25, Chapter 4])

$$(A \otimes B)^\top = A^\top \otimes B^\top, \tag{13}$$

$$(A \otimes B)(C \otimes D) = (AC) \otimes (BD). \tag{14}$$

The following lemmas can be found in [47].

Lemma 3.1 *For any two vectors u, v and a symmetric matrix W, we have*

$$2u^\top Wv = \|u\|_W^2 + \|v\|_W^2 - \|u - v\|_W^2. \tag{15}$$

Lemma 3.2 *Given a function ϕ and a fixed point $\bar{x}$, if for any λ, it holds that*

$$F(\bar{x}) - F(x^*) - \langle \lambda, A\bar{x} - b\rangle \le \phi(\lambda),$$

then for any $\gamma > 0$, we have

$$F(\bar{x}) - F(x^*) + \gamma\|A\bar{x} - b\| \le \sup_{\|\lambda\| \le \gamma} \phi(\lambda).$$

Lemma 3.3 *For any $\epsilon \ge 0$, if $F(\bar{x}) - F(x^*) + \gamma\|A\bar{x} - b\| \le \epsilon$, then we have*

$$\|A\bar{x} - b\| \le \frac{\epsilon}{\gamma - \|\lambda^*\|} \text{ and } -\frac{\|\lambda^*\|\epsilon}{\gamma - \|\lambda^*\|} \le F(\bar{x}) - F(x^*) \le \epsilon,$$

where (x^*, λ^*) satisfies the KKT conditions in (11), and we assume $\|\lambda^*\| < \gamma$.

3.2 Technical assumptions

Throughout our analysis, we make the following assumptions.

Assumption 1 *The solution set $\mathcal{X}^*$ is nonempty, i.e., there exists a pair (x^*, λ^*) satisfying the KKT conditions in (11).*

Assumption 2 *The solution set of the subproblem in (7) is nonempty for all i.*

Assumption 3 *The mixing matrices $\hat{W}$ and $\tilde{W}$ satisfy:*

$$\hat{w}_{ij} = \tilde{w}_{ij} = 1, \forall j \ge i, \tag{16}$$

$$\exists\, \hat{u}, \tilde{u} \text{ such that } \hat{U} = \hat{W} - e\hat{u}^\top,\ \tilde{U} = \tilde{W} - e\tilde{u}^\top \text{ are symmetric}, \tag{17}$$

where e is the all-one vector.

The first assumption is minimal for finding a solution to (5). The second one is for well-definedness of the proposed algorithm, and it can be guaranteed if P_i's are all positive definite.

The requirements in (16) are for easy implementation of the update in (7) because otherwise x_i^{k+1} may implicitly depend on the later updated block variables. The conditions in (17) are for technical reason; see (20) below. They can be satisfied by first choosing $\hat{u}, \tilde{u}$ and then determining the corresponding $\hat{W}, \tilde{W}$; see the formula (37) below. How to choose $\hat{u}, \tilde{u}$ will be discussed in the next section since it is inspired from our convergence analysis.

3.3 Convergence results of Algorithm 1

We show that with appropriate proximal terms, Algorithm 1 can have $O(1/t)$ convergence rate, where t is the number of iterations. The result includes several existing ones as special cases, and we will discuss it after presenting our convergence result.

We first establish a few inequalities.

Proposition 3.4 *Let $\hat{W}, \tilde{W}$ satisfy the conditions in Assumption 3 and define*

$$y_i = H_i x_i, \quad z_i = A_i x_i, \, \forall i.$$

Then for any $\alpha_1, \alpha_2 > 0$,

$$\begin{aligned}
&\sum_{i=1}^m \sum_{j=1}^m \hat{w}_{ij} \langle H_i(x_i^{k+1} - x_i), H_j(x_j^{k+1} - x_j^k) \rangle \qquad (18)\\
\leq &\frac{1}{2}\left(\|y^{k+1} - y\|_{\hat{V}}^2 - \|y^k - y\|_{\hat{V}}^2 + \|y^{k+1} - y^k\|_{\hat{V}}^2 \right) + \frac{1}{2\alpha_1} \|H(x^{k+1} - x)\|^2 + \frac{\alpha_1}{2} \left\| \left(\hat{u}^\top \otimes I\right)(y^{k+1} - y^k) \right\|^2,
\end{aligned}$$

and

$$\begin{aligned}
&\sum_{i=1}^m \sum_{j=1}^m \tilde{w}_{ij} \langle A_i(x_i^{k+1} - x_i), A_j(x_j^{k+1} - x_j^k) \rangle \qquad (19)\\
\leq &\frac{1}{2}\left(\|z^{k+1} - z\|_{\tilde{V}}^2 - \|z^k - z\|_{\tilde{V}}^2 + \|z^{k+1} - z^k\|_{\tilde{V}}^2 \right) + \frac{1}{2\alpha_2} \|A(x^{k+1} - b)\|^2 + \frac{\alpha_2}{2} \left\| \left(\tilde{u}^\top \otimes I\right)(z^{k+1} - z^k) \right\|^2,
\end{aligned}$$

where $\hat{V} = \hat{U} \otimes I$ and $\tilde{V} = \tilde{U} \otimes I$.

Proof. We only show (18), and (19) follows in the same way. By the definition of y, we have

$$\begin{aligned}
&\sum_{i=1}^m \sum_{j=1}^m \hat{w}_{ij} \langle H_i(x_i^{k+1} - x_i), H_j(x_j^{k+1} - x_j^k) \rangle\\
=&(y^{k+1} - y)^\top \left(\hat{W} \otimes I\right)(y^{k+1} - y^k)\\
\overset{(17)}{=}&(y^{k+1} - y)^\top \left(\hat{U} \otimes I\right)(y^{k+1} - y^k) + (y^{k+1} - y)^\top \left((e\hat{u}^\top) \otimes I\right)(y^{k+1} - y^k)\\
\overset{(14)}{=}&(y^{k+1} - y)^\top \left(\hat{U} \otimes I\right)(y^{k+1} - y^k) + (y^{k+1} - y)^\top \left((e \otimes I)(\hat{u}^\top \otimes I)\right)(y^{k+1} - y^k)\\
=&(y^{k+1} - y)^\top \left(\hat{U} \otimes I\right)(y^{k+1} - y^k) + \left(H(x^{k+1} - x)\right)^\top \left(\hat{u}^\top \otimes I\right)(y^{k+1} - y^k) \qquad (20)\\
\leq&(y^{k+1} - y)^\top \left(\hat{U} \otimes I\right)(y^{k+1} - y^k) + \frac{1}{2\alpha_1} \|H(x^{k+1} - x)\|^2 + \frac{\alpha_1}{2} \left\| \left(\hat{u}^\top \otimes I\right)(y^{k+1} - y^k) \right\|^2\\
\overset{(15)}{=}&\frac{1}{2}\left(\|y^{k+1} - y\|_{\hat{V}}^2 - \|y^k - y\|_{\hat{V}}^2 + \|y^{k+1} - y^k\|_{\hat{V}}^2 \right) + \frac{1}{2\alpha_1} \|H(x^{k+1} - x)\|^2 + \frac{\alpha_1}{2} \left\| \left(\hat{u}^\top \otimes I\right)(y^{k+1} - y^k) \right\|^2,
\end{aligned}$$

where the inequality follows from the Cauchy-Schwartz inequality. □

Proposition 3.5 (One-iteration result) *Let $\{(x^k, \lambda^k)\}_{k\geq 1}$ be the sequence generated from Algorithm 1. Under Assumptions 2 and 3, for any (x, λ) such that $Ax = b$, it holds that*

$$\begin{aligned}
&F(x^{k+1}) - F(x) - \langle \lambda, Ax^{k+1} - b\rangle + \frac{1}{2\rho}\left(\|\lambda^{k+1} - \lambda\|^2 - \|\lambda^k - \lambda\|^2\right) + \frac{\alpha_1 - 1}{2\alpha_1}\|H(x^{k+1} - x)\|^2 \\
&+ \left(\frac{\beta - \rho}{2} + \frac{\beta(\alpha_2 - 1)}{2\alpha_2}\right)\|Ax^{k+1} - b\|^2 + \frac{1}{2}\big(\|x^{k+1} - x\|_P^2 - \|x^k - x\|_P^2 + \|x^{k+1} - x^k\|_P^2\big) \\
\leq &\frac{1}{2}\left(\|y^{k+1} - y\|_{\hat{V}}^2 - \|y^k - y\|_{\hat{V}}^2 + \|y^{k+1} - y^k\|_{\hat{V}}^2\right) + \frac{\alpha_1}{2}\left\|\big(\hat{u}^\top \otimes I\big)(y^{k+1} - y^k)\right\|^2 \\
&+ \frac{\beta}{2}\left(\|z^{k+1} - z\|_{\tilde{V}}^2 - \|z^k - z\|_{\tilde{V}}^2 + \|z^{k+1} - z^k\|_{\tilde{V}}^2\right) + \frac{\beta\alpha_2}{2}\left\|\big(\tilde{u}^\top \otimes I\big)(z^{k+1} - z^k)\right\|^2,
\end{aligned} \tag{21}$$

where

$$P = \mathrm{blkdiag}(P_1, \ldots, P_m).$$

Proof. From the update (7), we have the optimality conditions: for $i = 1, \ldots, m$,

$$\begin{aligned}
0 =&\nabla_i f(\hat{x}^{k,i}) - A_i^\top \lambda^k + \tilde{\nabla} g_i(x_i^{k+1}) + \beta A_i^\top\big(A\tilde{x}^{k,i} - b\big) + P_i(x_i^{k+1} - x_i^k) \\
=&H_i^\top\left(Hx^{k+1} - \sum_{j=1}^m \hat{w}_{ij} H_j(x_j^{k+1} - x_j^k)\right) - A_i^\top \lambda^k + \tilde{\nabla} g_i(x_i^{k+1}) \\
&+ \beta A_i^\top\left((Ax^{k+1} - b) - \sum_{j=1}^m \tilde{w}_{ij} A_j(x_j^{k+1} - x_j^k)\right) + P_i(x_i^{k+1} - x_i^k),
\end{aligned} \tag{22}$$

where $\tilde{\nabla} g_i(x_i^{k+1})$ is a subgradient of g_i at x_i^{k+1}. Doing inner product of both sides of (22) with $x_i^{k+1} - x_i$, and summing them together over i, we have

$$\begin{aligned}
&\sum_{i=1}^m\sum_{j=1}^m \hat{w}_{ij}\big\langle H_i(x_i^{k+1} - x_i), H_j(x_j^{k+1} - x_j^k)\big\rangle + \beta\sum_{i=1}^m\sum_{j=1}^m \tilde{w}_{ij}\big\langle A_i(x_i^{k+1} - x_i), A_j(x_j^{k+1} - x_j^k)\big\rangle \\
=&\big\langle H(x^{k+1} - x), Hx^{k+1}\big\rangle + \big\langle A(x^{k+1} - x), -\lambda^k + \beta(Ax^{k+1} - b)\big\rangle + \big\langle x^{k+1} - x, \tilde{\nabla} g(x^{k+1})\big\rangle \\
&+ \big\langle x^{k+1} - x, P(x^{k+1} - x^k)\big\rangle \\
\geq&\big\langle H(x^{k+1} - x), Hx^{k+1}\big\rangle + \big\langle A(x^{k+1} - x), -\lambda^k + \beta(Ax^{k+1} - b)\big\rangle + g(x^{k+1}) - g(x) \\
&+ \big\langle x^{k+1} - x, P(x^{k+1} - x^k)\big\rangle \\
=&\frac{1}{2}\big(\|H(x^{k+1} - x)\|^2 - \|Hx\|^2 + \|Hx^{k+1}\|^2\big) - \langle \lambda^{k+1}, Ax^{k+1} - b\rangle \\
&+ (\beta - \rho)\|Ax^{k+1} - b\|^2 + g(x^{k+1}) - g(x) + \frac{1}{2}\big(\|x^{k+1} - x\|_P^2 - \|x^k - x\|_P^2 + \|x^{k+1} - x^k\|_P^2\big) \\
=&\frac{1}{2}\|H(x^{k+1} - x)\|^2 + F(x^{k+1}) - F(x) - \langle \lambda^{k+1}, Ax^{k+1} - b\rangle \\
&+ (\beta - \rho)\|Ax^{k+1} - b\|^2 + \frac{1}{2}\big(\|x^{k+1} - x\|_P^2 - \|x^k - x\|_P^2 + \|x^{k+1} - x^k\|_P^2\big),
\end{aligned} \tag{23}$$

where the inequality uses the convexity of g, and in the second equality, we have used (15), the update rule (8), and the condition $Ax = b$.

Substituting (18) and (19) into (23), we have

$$\begin{aligned}
&F(x^{k+1}) - F(x) - \langle \lambda^{k+1}, Ax^{k+1} - b\rangle + \frac{\alpha_1 - 1}{2\alpha_1}\|H(x^{k+1} - x)\|^2 \\
&+ \left(\frac{\beta}{2} + \frac{\beta(\alpha_2 - 1)}{2\alpha_2} - \rho\right)\|Ax^{k+1} - b\|^2 + \frac{1}{2}\big(\|x^{k+1} - x\|_P^2 - \|x^k - x\|_P^2 + \|x^{k+1} - x^k\|_P^2\big) \\
\le &\frac{1}{2}\left(\|y^{k+1} - y\|_{\hat{V}}^2 - \|y^k - y\|_{\hat{V}}^2 + \|y^{k+1} - y^k\|_{\hat{V}}^2\right) + \frac{\alpha_1}{2}\left\|(\hat{u}^\top \otimes I)(y^{k+1} - y^k)\right\|^2 \\
&+ \frac{\beta}{2}\left(\|z^{k+1} - z\|_{\tilde{V}}^2 - \|z^k - z\|_{\tilde{V}}^2 + \|z^{k+1} - z^k\|_{\tilde{V}}^2\right) + \frac{\beta\alpha_2}{2}\left\|(\tilde{u}^\top \otimes I)(z^{k+1} - z^k)\right\|^2. \qquad (24)
\end{aligned}$$

From the update (8), we have

$$\begin{aligned}
0 =&\langle \lambda^{k+1} - \lambda, Ax^{k+1} - b\rangle + \frac{1}{\rho}\langle \lambda^{k+1} - \lambda, \lambda^{k+1} - \lambda^k\rangle \\
\overset{(15)}{=}&\langle \lambda^{k+1} - \lambda, Ax^{k+1} - b\rangle + \frac{1}{2\rho}\left(\|\lambda^{k+1} - \lambda\|^2 - \|\lambda^k - \lambda\|^2 + \|\lambda^{k+1} - \lambda^k\|^2\right) \\
=&\langle \lambda^{k+1} - \lambda, Ax^{k+1} - b\rangle + \frac{1}{2\rho}\left(\|\lambda^{k+1} - \lambda\|^2 - \|\lambda^k - \lambda\|^2\right) + \frac{\rho}{2}\|A(x^{k+1} - b)\|^2. \qquad (25)
\end{aligned}$$

Summing (24) and (25) together gives the desired result. □

Now we are ready to present our main result.

Theorem 3.6 *Under Assumptions 1 through 3, let $\{(x^k, \lambda^k)\}_{k\ge 1}$ be the sequence generated from Algorithm 1 with parameters:*

$$\beta \ge \rho > 0, \qquad (26)$$

$$P - D_H^\top\big((\hat{W} - e\hat{u}^\top + \alpha_1\hat{u}\hat{u}^\top) \otimes I\big)D_H - \beta D_A^\top\big((\tilde{W} - e\tilde{u}^\top + \alpha_2\tilde{u}\tilde{u}^\top) \otimes I\big)D_A := \hat{P} \succeq 0. \qquad (27)$$

where $\alpha_1, \alpha_2 \ge 1$, and

$$D_H = \mathrm{blkdiag}(H_1, \ldots, H_m), \quad D_A = \mathrm{blkdiag}(A_1, \ldots, A_m).$$

Let $\bar{x}^{k+1} = \sum_{k=1}^t \frac{x^{k+1}}{t}$. Then

$$\big|F(\bar{x}^{t+1}) - F(x^*)\big| \le \frac{1}{2t}\left(\frac{4}{\rho}\|\lambda^*\|^2 + \|x^1 - x^*\|_P^2 - \|y^1 - y^*\|_{\hat{V}}^2 - \beta\|z^1 - z^*\|_{\tilde{V}}^2\right), \qquad (28a)$$

$$\|A\bar{x}^{t+1} - b\| \le \frac{1}{2t\|\lambda^*\|}\left(\frac{4}{\rho}\|\lambda^*\|^2 + \|x^1 - x^*\|_P^2 - \|y^1 - y^*\|_{\hat{V}}^2 - \beta\|z^1 - z^*\|_{\tilde{V}}^2\right), \qquad (28b)$$

where $\hat{V}$ and $\tilde{V}$ are defined in Proposition 3.4, and (x^*, λ^*) is any point satisfying the KKT conditions in (11).

In addition, if $\alpha_2 > 1$ and $\hat{P} \succ 0$, then (x^k, λ^k) converges to a point $(x^\infty, \lambda^\infty)$ that satisfies the KKT conditions in (11).

Proof. Summing the inequality (21) from $k=1$ through t and noting $\beta \geq \rho$, we have

$$\begin{aligned}
&\sum_{k=1}^{t}\left[F(x^{k+1})-F(x)-\langle\lambda, Ax^{k+1}-b\rangle\right]+\frac{1}{2\rho}\|\lambda^{t+1}-\lambda\|^2+\frac{1}{2}\|x^{t+1}-x\|_P^2\\
&+\sum_{k=1}^{t}\left(\frac{1}{2}\|x^{k+1}-x^k\|_P^2+\frac{\alpha_1-1}{2\alpha_1}\|H(x^{k+1}-x)\|^2+\frac{\beta(\alpha_2-1)}{2\alpha_2}\|Ax^{k+1}-b\|^2\right)\\
\leq&\frac{1}{2\rho}\|\lambda^1-\lambda\|^2+\frac{1}{2}\|x^1-x\|_P^2-\frac{1}{2}\|y^1-y\|_{\hat{V}}^2-\frac{\beta}{2}\|z^1-z\|_{\tilde{V}}^2\\
&+\frac{1}{2}\|y^{t+1}-y\|_{\hat{V}}^2+\frac{1}{2}\sum_{k=1}^{t}\|y^{k+1}-y^k\|_{\hat{V}}^2+\frac{\alpha_1}{2}\sum_{k=1}^{t}\left\|\left(\hat{u}^\top\otimes I\right)(y^{k+1}-y^k)\right\|^2\\
&+\frac{\beta}{2}\|z^{t+1}-z\|_{\tilde{V}}^2+\frac{\beta}{2}\sum_{k=1}^{t}\|z^{k+1}-z^k\|_{\tilde{V}}^2+\frac{\beta\alpha_2}{2}\sum_{k=1}^{t}\left\|\left(\tilde{u}^\top\otimes I\right)(z^{k+1}-z^k)\right\|^2.
\end{aligned}\tag{29}$$

Note that

$$\begin{aligned}
&\|y^{k+1}-y^k\|_{\hat{V}}^2+\alpha_1\left\|\left(\hat{u}^\top\otimes I\right)(y^{k+1}-y^k)\right\|^2\\
=&(y^{k+1}-y^k)^\top(\hat{V}+\alpha_1\hat{u}\hat{u}^\top\otimes I)(y^{k+1}-y^k)\\
=&(x^{k+1}-x^k)^\top D_H^\top\left((\hat{W}-e\hat{u}^\top+\alpha_1\hat{u}\hat{u}^\top)\otimes I\right)D_H(x^{k+1}-x^k),
\end{aligned}$$

and similarly

$$\begin{aligned}
&\|z^{k+1}-z^k\|_{\tilde{V}}^2+\alpha_2\left\|\left(\tilde{u}^\top\otimes I\right)(z^{k+1}-z^k)\right\|^2\\
=&(x^{k+1}-x^k)^\top D_A^\top\left((\tilde{W}-e\tilde{u}^\top+\alpha_2\tilde{u}\tilde{u}^\top)\otimes I\right)D_A(x^{k+1}-x^k).
\end{aligned}$$

Hence, by the choice of P in (27), we have from (29) that

$$\begin{aligned}
&\sum_{k=1}^{t}\left[F(x^{k+1})-F(x)-\langle\lambda, Ax^{k+1}-b\rangle\right]+\frac{1}{2\rho}\|\lambda^{t+1}-\lambda\|^2\\
&+\sum_{k=1}^{t}\left(\frac{1}{2}\|x^{k+1}-x^k\|_{\hat{P}}^2+\frac{\alpha_1-1}{2\alpha_1}\|H(x^{k+1}-x)\|^2+\frac{\beta(\alpha_2-1)}{2\alpha_2}\|Ax^{k+1}-b\|^2\right)\\
\leq&\frac{1}{2\rho}\|\lambda^1-\lambda\|^2+\frac{1}{2}\|x^1-x\|_P^2-\frac{1}{2}\|y^1-y\|_{\hat{V}}^2-\frac{\beta}{2}\|z^1-z\|_{\tilde{V}}^2.
\end{aligned}\tag{30}$$

Since $\alpha_1, \alpha_2 \geq 1$ and $\hat{P} \succeq 0$, it follows from the above inequality and the convexity of F that

$$\begin{aligned}
&F(\bar{x}^{t+1})-F(x^*)-\langle\lambda, A\bar{x}^{t+1}-b\rangle\\
\leq&\frac{1}{2t}\left(\frac{1}{\rho}\|\lambda^1-\lambda\|^2+\|x^1-x^*\|_P^2-\|y^1-y^*\|_{\hat{V}}^2-\beta\|z^1-z^*\|_{\tilde{V}}^2\right).
\end{aligned}\tag{31}$$

Since $\lambda^1=0$, we use Lemmas 3.2 and 3.3 with $\gamma=2\|\lambda^*\|$ to have (28).

If $\alpha_2 > 1$ and $\hat{P} \succ 0$, then letting $\lambda = \lambda^*$ in (30) and also using (12), we have

$$\lim_{k\to\infty}(x^{k+1} - x^k) = 0, \quad \lim_{k\to\infty}(\lambda^{k+1} - \lambda^k) = -\frac{1}{\rho}\lim_{k\to\infty}(Ax^{k+1} - b) = 0, \tag{32}$$

and thus

$$\lim_{k\to\infty}(x^k - \tilde{x}^{k,i}) = \lim_{k\to\infty}(x^k - \hat{x}^{k,i}) = 0, \ \forall i. \tag{33}$$

On the other hand, letting $(x, \lambda) = (x^*, \lambda^*)$ in (21), using (12), and noting $\alpha_1, \alpha_2 \geq 1$, we have

$$\begin{aligned}&\frac{1}{2\rho}\|\lambda^{k+1} - \lambda^*\|^2 + \frac{1}{2}\left(\|x^{k+1} - x^*\|_P^2 - \|y^{k+1} - y^*\|_{\hat{V}}^2 - \beta\|z^{k+1} - z^*\|_{\tilde{V}}^2\right)\\ \leq&\frac{1}{2\rho}\|\lambda^k - \lambda^*\|^2 + \frac{1}{2}\left(\|x^k - x^*\|_P^2 - \|y^k - y^*\|_{\hat{V}}^2 - \beta\|z^k - z^*\|_{\tilde{V}}^2\right),\end{aligned} \tag{34}$$

which together with the choice of P indicates the boundedness of $\{(x^k, \lambda^k)\}_{k\geq 1}$. Hence, it must have a finite limit point $(x^\infty, \lambda^\infty)$, and there is a subsequence $\{(x^k, \lambda^k)\}_{k\in\mathcal{K}}$ convergent to this limit point. From (32), it immediately follows that $Ax^\infty - b = 0$. In addition, letting $\mathcal{K} \ni k \to \infty$ in (7) and using (32) and (33) gives

$$x_i^\infty = \arg\min_{x_i}\langle \nabla_i f(x^\infty) - A_i^\top \lambda^\infty, x_i\rangle + g_i(x_i) + \frac{1}{2}\|x_i - x_i^\infty\|_{P_i}^2, \ \forall i,$$

and thus the optimality condition holds:

$$0 \in \nabla_i f(x^\infty) + \partial g_i(x_i^\infty) - A_i^\top \lambda^\infty, \ \forall i.$$

Therefore, $(x^\infty, \lambda^\infty)$ satisfies the conditions in (11). Since (34) holds for any point (x^*, λ^*) satisfying (11), it also holds with $(x^*, \lambda^*) = (x^\infty, \lambda^\infty)$. Because $(x^\infty, \lambda^\infty)$ is a limit point of $\{(x^k, \lambda^k)\}_{k\geq 1}$, we obtain the convergence of (x^k, λ^k) to $(x^\infty, \lambda^\infty)$ and complete the proof. □

4 How to choose mixing matrices

In this section, we discuss how to choose $\hat{W}$ and $\tilde{W}$ such that they satisfy Assumption 3. Note that both $\hat{W}$ and $\tilde{W}$ satisfy:

$$w_{ij} = 1, \ \forall j \geq i, \tag{35}$$

$$\exists\, u \text{ such that } W - eu^\top = W^\top - ue^\top. \tag{36}$$

Hence, their upper triangular parts have been fixed, and we only need to set their strictly lower triangular parts. Denote $\mathcal{U}(W)$ and $\mathcal{L}(W)$ respectively as the upper and strictly lower triangular parts of W, i.e., $W = \mathcal{L}(W) + \mathcal{U}(W)$, and thus (36) is equivalent to requiring the existence of u such that

$$\mathcal{L}(W) + \mathcal{U}(W) - eu^\top = \mathcal{L}(W)^\top + \mathcal{U}(W)^\top - ue^\top.$$

It suffices to let

$$\mathcal{L}(W) = \mathcal{L}(W^\top + eu^\top - ue^\top) = \mathcal{L}(eu^\top - ue^\top + E). \tag{37}$$

Therefore, given any vector u, we can find a corresponding W by setting its upper triangular part to all one's and its strictly lower triangular part according to the above formula.

4.1 Finding u by solving SDP

Theoretically, proximal terms help convergence guarantee of the algorithm. However, empirically, these terms can slow the convergence speed. Based on these observations, we aim at finding a block diagonal P such that (27) holds and also is as close to *zero* as possible.

Let $\hat{D} = \mathrm{diag}(\hat{D}_1, \ldots, \hat{D}_m)$ with each $\hat{D}_i \in \{0, 1\}$, where *one* means linearization to f at x_i^k and *zero* no linearization; see (10). Similarly, we let $\tilde{D} = \mathrm{diag}(\tilde{D}_1, \ldots, \tilde{D}_m)$ with each $\tilde{D}_i \in \{0, 1\}$. Then one choice of the block diagonal matrix P satisfying (27) could be

$$P_i = (1 - \hat{D}_i)H_i^\top H_i + \beta(1 - \tilde{D}_i)A_i^\top A_i + d_i(\|H_i\|_2^2 + \beta\|A_i\|_2^2)I,\ i = 1, \ldots m. \tag{38}$$

where

$$\mathrm{diag}(d_1, \ldots, d_m) \succeq (\hat{W} - e\hat{u}^\top + \alpha_1\hat{u}\hat{u}^\top) - (I - \hat{D}) + \beta(\tilde{W} - e\tilde{u}^\top + \alpha_2\tilde{u}\tilde{u}^\top) - \beta(I - \tilde{D}). \tag{39}$$

To obtain P according to (38), we only need to specify the values of d_i's. Recall that we aim at finding a close-to-zero P, so it would be desirable to choose vectors $\hat{u}$ and $\tilde{u}$ such that $\max_i d_i$ is as small as possible. Without assuming specific structures to H or A, we simply set $\alpha_1 = \alpha_2 = 1$, $\hat{u} = \tilde{u} = u$ and $\hat{D} = \tilde{D} = D$, and from the formula (37), it follows $\hat{W} = \tilde{W} = W$. Therefore, to minimize $\max_i d_i$, we solve the following optimization problem:

$$\min_{u,W} \lambda_{\max}\left(W - eu^\top + uu^\top - I + D\right),\ \text{s.t. } W - eu^\top = W^\top - ue^\top, \tag{40}$$

where $\lambda_{\max}(B)$ denotes the maximum eigenvalue of a symmetric matrix B.

Using (37), we represent W by u and write (40) equivalently to

$$\min_{u} \lambda_{\max}\left(\mathcal{L}(eu^\top - ue^\top) + E - eu^\top + uu^\top - I + D\right), \tag{41}$$

which can be further formulated as an SDP by the relation between the positive-definiteness of a 2×2 block matrix and its Schur complement (c.f., [1, Appendix A.5.5]):

$$\begin{bmatrix} A & B \\ B^\top & C \end{bmatrix} \succeq 0 \Leftrightarrow A - BC^{-1}B^\top \succeq 0, \tag{42}$$

where A is symmetric and $C \succ 0$. Let

$$\sigma = \lambda_{\max}\left(\mathcal{L}(eu^\top - ue^\top) + E - eu^\top + uu^\top - I + D\right).$$

Then

$$\sigma I \succeq \mathcal{L}(eu^\top - ue^\top) + E - eu^\top + uu^\top - I + D,$$

and by (42) it is equivalent to

$$\begin{bmatrix} (\sigma + 1)I - D - \mathcal{L}(eu^\top - ue^\top) - E + eu^\top & u \\ u^\top & 1 \end{bmatrix} \succeq 0.$$

Therefore, (41) is equivalent to the SDP:

$$\min_{\sigma,u} \ \sigma, \ \text{s.t.} \ \begin{bmatrix} (\sigma+1)I - D - \mathcal{L}(eu^\top - ue^\top) - E + eu^\top & u \\ u^\top & 1 \end{bmatrix} \succeq 0. \tag{43}$$

Note that the problem (5) can be extremely large. However, the SDP (43) could be much smaller and can be efficiently and accurately solved by the interior-point method. In addition, (43) does not depend on the data matrix H and A, so we can solve it offline. We use SDPT3 [39] to find its solution in our numerical tests.

If the sizes of H and A are not large, upon solving (43), one can explicitly form the matrix $D_H^\top\big((W - eu^\top + uu^\top) \otimes I\big)D_H + \beta D_A^\top\big((W - eu^\top + uu^\top) \otimes I\big)D_A$ and compute its spectral norm. This way, one can have a smaller P. However, for large-scale H or A, it can be overwhelmingly expensive to do so.

In addition, note that we can add more constraints to (43) to obtain a desired W. For instance, we can partition the m blocks into several groups. Then we update the blocks in the same group in parallel in Jacobian manner and cyclically renew the groups. This corresponds to fix a few block matrices in the lower triangular part of W to all one's.

4.2 Special cases

A few special cases are as follows.

- If u is the *zero* vector, then by (37), we have the lower triangular part of W to be all one's, and this way gives a fully Jacobian BCU method. Hence, Theorem 3.6 applies for the Jacobian update method. In this case, if $\hat{D} = \tilde{D} = 0$ in (39), then we have $\min \max_i d_i \geq (1+\beta)m, \forall i$ that is significantly greater than the optimal value of (43).

- If we enforce the lower triangular part of W to all *zero*'s, Algorithm 1 will reduce to a fully Gauss-Seidel BCU method. However, adding such constraints into (43) would lead to infeasibility. Hence, Theorem 3.6 does not apply to this case.

- If $m = 2$ and $D = 0$, i.e., there are only two blocks and no linearization is performed, solving (43) would give the solution $\sigma = 0$ and $u = (0,1)^\top$. This way, we have

$$W = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}, \quad P = \begin{bmatrix} H_1^\top H_1 + \beta A_1^\top A_1 & 0 \\ 0 & H_2^\top H_2 + \beta A_2^\top A_2 \end{bmatrix}$$

  and thus recover the 2-block ADMM with nonseparable quadratic term f in the objective. Theorem 3.6 implies $O(1/t)$ convergence rate for this case, and it improves the result in [5], which shows convergence of this special case but without rate estimate.

5 Numerical experiments

In this section, we apply Algorithm 1 to three problems: quadratic programming, compressive principal component pursuit, and the multi-class support vector machine problem. We test it with two different mixing matrices: all-one matrix and the one given by the method discussed in section 4.1. The former corresponds to a fully Jacobian update method and the latter to a hybrid Jacobian and Gauss-Seidel method, dubbed as Jacobi-PC and JaGS-PC respectively. We compare them to a recently proposed randomized proximal block coordinate update method (named as random-PC) in [14], the ADMM with Gauss back substitution (named as ADMM-GBS) in [19], and also the direct ADMM. Note that the direct ADMM is not guaranteed to converge for problems with more than two blocks, but empirically it can often perform well. ADMM-GBS is designed for separable multi-block convex problems. It does not allow linearization to the augmented term. In addition, it requres all A_i's to be full-column rank. Hence, the proposed algorithm is applicable to broader class of problems, but we observe that JaGS-PC can be comparable to direct ADMM and ADMM-GBS.

5.1 Adaptive proximal terms

As we mentioned previously, the proximal terms used in (7) help the convergence guarantee but can empirically slow the convergence speed. Here, we set P_i's similar to (38) but with a simple adaptive way as follows:

$$P_i^k = (1-\hat{D}_i)H_i^\top H_i + \beta(1-\tilde{D}_i)A_i^\top A_i + d^k(\|H_i\|_2^2 + \beta\|A_i\|_2^2)I,\ i = 1,\dots m, \tag{44}$$

where we set $\hat{D} = \tilde{D} = D$ as discussed in section 4.1. After each iteration k, we check if the following inequality holds

$$\begin{aligned}\eta\|x^{k+1}-x^k\|_{P^k}^2 \le &\|y^{k+1}-y^k\|_{\hat{V}}^2 + \left\|(\hat{u}^\top\otimes I)(y^{k+1}-y^k)\right\|^2 \\ &+ \beta\|z^{k+1}-z^k\|_{\tilde{V}}^2 + \beta\left\|(\tilde{u}^\top\otimes I)(z^{k+1}-z^k)\right\|^2,\end{aligned} \tag{45}$$

and set

$$d^k = \begin{cases} \min\left(d^{k-1}+d_{\text{inc}}, d_{\max}\right), & \text{if (45) holds} \\ d^{k-1}, & \text{otherwise}\end{cases} \tag{46}$$

where we set $\hat{u} = \tilde{u} = u$, d_{inc} is a small positive number, and $\eta = 0.999$ is used in all the tests. For Jacobi-PC, we set $d_{\max} = \lambda_{\max}(W - I + D)$ with the same W and D as those in (40), and for JaGS-PC, we set $d_{\max}$ to the optimal value of (43). Note that d^k can only be incremented in finitely many times, and thus (45) can only happen in finitely many iterations. In addition, note that both sides of (45) can be evaluated as cheaply as computing $x^\top Qx$ and $x^\top A^\top Ax$.

The above adaptive way of setting P^k is inspired from our convergence analysis. From (29), we see that if y^{t+1} and z^{t+1} are bounded and (45) only happens in finitely many times, the sublinear convergence result in Theorem 3.6 still holds.

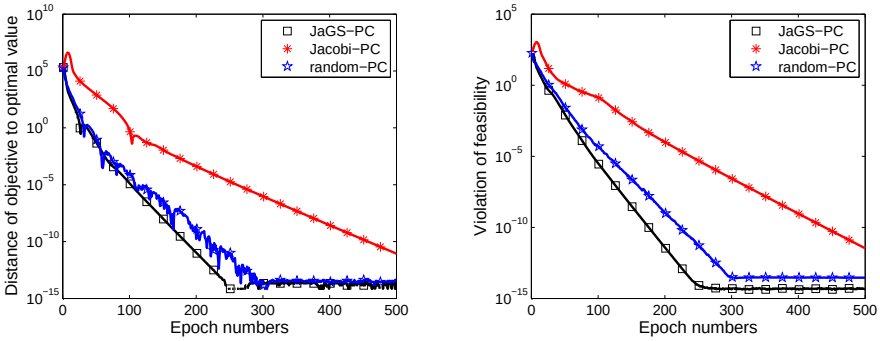


Figure 1: Results by applying three different proximal block coordinate update methods: Jacobi-PC, JaGS-PC, and random-PC on the quadratic programming (47). Left: the distance of the objective to optimal value $|F(x^k) - F(x^*)|$; Right: the violation of feasibility $\|Ax^k - b\|$.

5.2 Quadratic programming

We test Jacobi-PC and JaGS-PC on the nonnegative linearly constrained quadratic programming

$$\min_x F(x) = \frac{1}{2}x^\top Qx + c^\top x, \text{ s.t. } Ax = b,\ x \geq 0, \tag{47}$$

where $Q \in \mathbb{R}^{n\times n}$ is a symmetric PSD matrix, $A \in \mathbb{R}^{p\times n}$ and $b \in \mathbb{R}^p, c \in \mathbb{R}^n$. In the test, we set $p = 200, n = 2000$ and $Q = H^\top H$ with $H \in \mathbb{R}^{(n-10)\times n}$ generated according to the standard Gaussian distribution. This generated Q is degenerate, and thus the problem is only weakly convex. The entries of c follow i.i.d. standard Gaussian distribution and those of b from uniform distribution on $[0, 1]$. We set $A = [B, I]$ to guarantee feasibility of the problem with B generated according to standard Gaussian distribution.

We evenly partition the variable x into $m = 40$ blocks, each one consisting of 50 coordinates. The same values of parameters are set for both Jacobi-PC and JaGS-PC as follows:

$$\beta = \rho = 1,\ \hat{D} = \tilde{D} = I,\ d^1 = 0.5,\ d_{\text{inc}} = 0.1.$$

They are compared to random-PC that uses the same penalty parameter $\beta = 1$ and $\rho = \frac{1}{m}$ according to the analysis in [14]. All three methods run to 500 epochs, where each epoch is equivalent to updating all blocks one time. Their per-epoch complexity is almost the same. To evaluate their performance, we compute the distance of the objective to optimal value $|F(x^k) - F(x^*)|$ and the violation of feasibility $\|Ax^k - b\|$ at each iteration k. Figure 1 plots the results by the three methods. From the figure, we see that JaGS-PC is significantly faster than Jacobi-PC in terms of both objective and feasibility, and random-PC is slightly slower than JaGS-PC.

5.3 Compressive principal component pursuit

In this subsection, we test Jacobi-PC and JaGS-PC on

$$\min_{X,Y} F(X,Y) = \mu\|X\|_1 + \|Y\|_*, \text{ s.t. } \mathcal{A}(X+Y) = b, \tag{48}$$

where $\|X\|_1 = \sum_{i,j}|X_{ij}|$, $\|Y\|_*$ denotes the matrix operator norm and equals the largest singular value of Y, $\mathcal{A}$ is a linear operator, and b contains the measurements. If $\mathcal{A}$ is the identity operator, (48) is called the principal component pursuit (PCP) proposed in [3], and it is compressive PCP [45] when $\mathcal{A}$ is an underdetermined measuring operator. We consider the sampling operator, i.e., $\mathcal{A} = \mathcal{P}_\Omega$, where Ω is an index set and $\mathcal{P}_\Omega$ is a projection keeping the entries in Ω and zeroing out all others.

Assume M to be the underlying matrix and $b = \mathcal{A}(M)$. Upon solving (48), $X+Y$ recovers M with sparse part X and low-rank part Y. We use the Escalator video dataset[1], which has 200 frames of 130×160 images. Each frame of 2D image is reshaped into a column vector to form a slim and tall matrix M with 200 columns. For this data, X will encode the foreground and Y the background of the video. We generate the index set Ω uniformly at random, and 30% samples are selected from each frame of image.

Introducing another variable Z, we write (48) equivalently to

$$\min_{X,Y,Z} \mu\|X\|_1 + \|Y\|_*, \text{ s.t. } X+Y=Z,\ \mathcal{A}(Z) = b, \tag{49}$$

which naturally has three block variables. Applying Algorithm 1 to (49) with adaptive P_i's as in (44) and $\tilde{D} = 0$ and noting $\|I + \mathcal{A}^\top\mathcal{A}\|_2 = 2$, we iteratively perform the updates:

$$X^{k+1} = \arg\min_X \mu\|X\|_1 - \langle\Lambda^k, X\rangle + \frac{\beta}{2}\|X+Y^k-Z^k\|_F^2 + \frac{d^k\beta}{2}\|X-X^k\|^2, \tag{50a}$$

$$Y^{k+1} = \arg\min_Y \|Y\|_* - \langle\Lambda^k, Y\rangle + \frac{\beta}{2}\|\tilde{X}^k+Y-Z^k\|_F^2 + \frac{d^k\beta}{2}\|Y-Y^k\|^2, \tag{50b}$$

$$Z^{k+1} = \arg\min_Z \langle\Lambda^k - \mathcal{A}^\top(\Pi^k), Z\rangle + \frac{\beta}{2}\|\tilde{X}^k+\tilde{Y}^k-Z^k\|_F^2 + \frac{\beta}{2}\|\mathcal{A}(Z)-b\|^2 + d^k\beta\|Z-Z^k\|^2, \tag{50c}$$

$$\Lambda^{k+1} = \Lambda^k - \rho(X^{k+1}+Y^{k+1}-Z^{k+1}), \tag{50d}$$

$$\Pi^{k+1} = \Pi^k - \rho(\mathcal{A}(Z^{k+1})-b). \tag{50e}$$

Since $\mathcal{A}\mathcal{A}^\top = I$, all three primal subproblems have closed-form solutions. We set $\beta = \rho = 0.05$ and $d_{\text{inc}} = 0.01$ for both JaGS-PC and Jacobi-PC methods, and $d^1 = 0$ for JaGS-PC and $d^1 = 1$ for Jacobi-PC because the latter can deviate from optimality very far away in the beginning if it starts with a small d^1. They are compared to the direct ADMM that sets $d^k = 0, \forall k$ in (50) and also ADMM-GBS that runs the direct ADMM first and then performs a correction step by Gauss

[1] Available from http://pages.cs.wisc.edu/~jiaxu/projects/gosus/supplement/

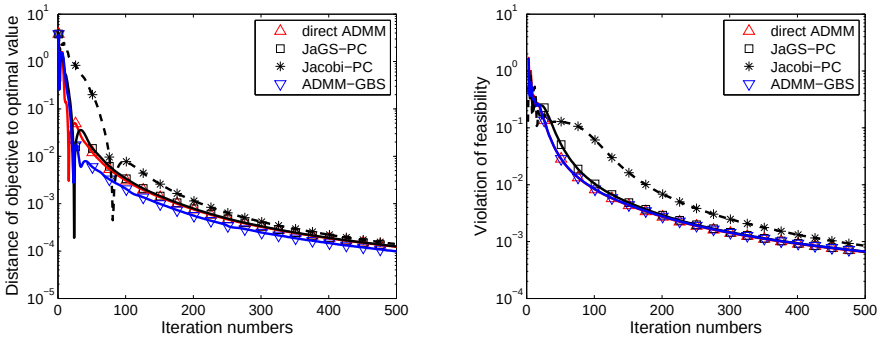


Figure 2: Results by four different methods for solving the compressive principal component pursuit (49) on the Escalator dataset. Left: the relative error between the objective and optimal value $\frac{|F(X^k,Y^k)-F(X^*,Y^*)|}{F(X^*,Y^*)}$; Right: relative violation of feasibility: $\frac{\|X^k+Y^k-Z^k\|_F+\|\mathcal{A}(Z^k)-b\|_F}{\|M\|_F}$.

back substitution. We use the same β and ρ for the direct ADMM and ADMM-GBS and set the correction step parameter of ADMM-GBS to 0.99. Figure 2 plots the results by all four methods. From the figure, we see that JaGS-PC performs significantly better than Jacobi-PC. JaGS-PC and direct ADMM perform almost the same, and ADMM-GBS is slightly better. Note that although direct ADMM works well on this example, its convergence is not guaranteed in general.

5.4 Multi-class support vector machine

In this subsection, we test Jacobi-PC and JaGS-PC on the multi-class support vector machine (MSVM) problem that is considered in [48]:

$$\min_X F(X) = \frac{1}{n}\sum_{i=1}^{n}\sum_{j=1}^{c} \iota_{\cdot\neq j}(b_i)[x_j^\top a_i + 1]_+ + \mu\|X\|_1, \text{ s.t. } Xe = 0, \tag{51}$$

where x_j is the j-th column of X, $\{(a_i, b_i)\}_{i=1}^n$ is the training dataset with label $b_i \in \{1, \ldots, c\}$, $\forall i$, $\iota_{\cdot\neq j}(b_i)$ equals *one* if $b_i \neq j$ and *zero* otherwise, and $[d]_+ = \max(0, d)$. We set the number of classes to $c = 3$ and randomly generate the data according to Gaussian distribution $\mathcal{N}(v_j, \Sigma_j)$ for the j-th class, where $v_j \in \mathbb{R}^p$ and $\Sigma_j \in \mathbb{R}^{p\times p}$ for $j = 1, 2, 3$ are

$$v_1 = \begin{bmatrix} E_{s\times 1} \\ 0 \end{bmatrix}, v_2 = \begin{bmatrix} 0_{s/2\times 1} \\ E_{s\times 1} \\ 0 \end{bmatrix}, v_3 = \begin{bmatrix} 0_{s\times 1} \\ E_{s\times 1} \\ 0 \end{bmatrix}, \Sigma_1 = \begin{bmatrix} \sigma E_{s\times s} + (1-\sigma)I & 0 \\ 0 & I \end{bmatrix},$$

$$\Sigma_2 = \begin{bmatrix} I_{\frac{s}{2}\times\frac{s}{2}} & 0 & 0 \\ 0 & \sigma E_{s\times s} + (1-\sigma)I & 0 \\ 0 & 0 & I \end{bmatrix}, \Sigma_3 = \begin{bmatrix} I_{s\times s} & 0 & 0 \\ 0 & \sigma E_{s\times s} + (1-\sigma)I & 0 \\ 0 & 0 & I \end{bmatrix},$$

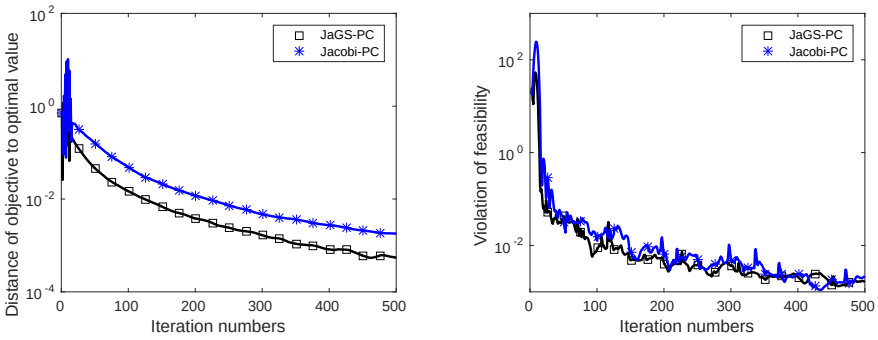


Figure 3: Results by Jacobi-PC and JaGS-PC for solving the multiclass support vector machine problem (51) on randomly generated data. Left: distance of the objective to optimal value $|F(X^k) - F(X^*)|$; Right: violation of feasibility $\|X^k e\|$.

where $E, I, 0$ respectively represent all-one, identity, and all-zero matrices of appropriate sizes, and the subscript specifies the size. The parameter σ measures correlation of features. This kind of dataset has also been used in [49] for testing binary SVM. In the test, we set $\sigma = 0.1, \mu = 0.001, p = 200$ and $n = 300$, each class consisting of 100 samples.

Letting $A = [a_1, \ldots, a_n]$ and $Y = A^\top X + 1$, we write (51) equivalently to

$$\min_{X,Y} \frac{1}{n} \sum_{i=1}^{n} \sum_{j=1}^{c} \iota_{\cdot \neq j}(b_i)[Y_{ij}]_+ + \mu \|X\|_1, \text{ s.t. } A^\top X - Y = 1, Xe = 0. \tag{52}$$

To apply Algorithm 1 to the above model, we partition the variable into four blocks (x_1, x_2, x_3, Y). Linearization to the augmented term is employed, i.e., $\tilde{D} = I$ in (44). The parameters are set to $\beta = \rho = 0.005$ and $d_{\text{inc}} = 0.1$ for both Jacobi-PC and JaGS-PC, and $d^1 = 0.5$ for JaGS-PC and $d^1 = 1$ for Jacobi-PC because again the latter can deviate from optimalty far away in the beginning if it starts with a small d^1. We plot the results in Figure 3, which shows that JaGS-PC performs significantly better than Jacobi-PC in terms of objective value.

6 Conclusions

We have proposed a hybrid Jacobian and Gauss-Seidel block coordinate update method for solving linearly constrained convex programming. The method performs each primal block variable update by minimizing a function that approximates the augmented Lagrangian at affinely combined points of the previous two iterates. We have presented a way to choose the mixing matrix with desired properties. Global iterate sequence convergence and also sublinear rate results of the hybrid method have been established. In addition, numerical experiments have been performed to demonstrate its efficiency.

References

[1] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004. 14

[2] X. Cai, D. Han, and X. Yuan. The direct extension of admm for three-block separable convex minimization models is convergent when one function is strongly convex. *Optimization Online*, 2014. 3

[3] E. J. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *Journal of the ACM (JACM)*, 58(3):11, 2011. 18

[4] C. Chen, B. He, Y. Ye, and X. Yuan. The direct extension of admm for multi-block convex minimization problems is not necessarily convergent. *Mathematical Programming*, 155(1-2):57–79, 2016. 3, 6

[5] C. Chen, M. Li, X. Liu, and Y. Ye. Extended admm and bcd for nonseparable convex minimization models with quadratic coupling terms: Convergence analysis and insights. *arXiv preprint arXiv:1508.00193*, 2015. 4, 15

[6] C. Chen, Y. Shen, and Y. You. On the convergence analysis of the alternating direction method of multipliers with three blocks. In *Abstract and Applied Analysis*, volume 2013. Hindawi Publishing Corporation, 2013. 3

[7] P. L. Combettes and J.-C. Pesquet. Stochastic quasi-fejér block-coordinate fixed point iterations with random sweeping. *SIAM Journal on Optimization*, 25(2):1221–1248, 2015. 3

[8] Y. Cui, X. Li, D. Sun, and K.-C. Toh. On the convergence properties of a majorized ADMM for linearly constrained convex optimization problems with coupled objective functions. *arXiv preprint arXiv:1502.00098*, 2015. 4

[9] C. Dang and G. Lan. Randomized methods for saddle point computation. *arXiv preprint arXiv:1409.8625*, 2014. 4

[10] D. Davis and W. Yin. A three-operator splitting scheme and its optimization applications. *arXiv preprint arXiv:1504.01032*, 2015. 3

[11] W. Deng, M.-J. Lai, Z. Peng, and W. Yin. Parallel multi-block admm with o (1/k) convergence. *arXiv preprint arXiv:1312.3040*, 2013. 3, 4

[12] J.-K. Feng, H.-B. Zhang, C.-Z. Cheng, and H.-M. Pei. Convergence analysis of l-admm for multi-block linear-constrained separable convex minimization problem. *Journal of the Operations Research Society of China*, 3(4):563–579, 2015. 6

[13] D. Gabay and B. Mercier. A dual algorithm for the solution of nonlinear variational problems via finite element approximation. *Computers & Mathematics with Applications*, 2(1):17–40, 1976. 3

[14] X. Gao, Y. Xu, and S. Zhang. Randomized primal-dual proximal block coordinate updates. *arXiv preprint arXiv:1605.05969*, 2016. 4, 5, 16, 17

[15] X. Gao and S. Zhang. First-order algorithms for convex optimization with nonseparate objective and coupled constraints. *Optimization online*, 3:5, 2015. 4

[16] R. Glowinski and A. Marrocco. Sur l'approximation, par eléments finis d'ordre un, et la résolution, par pénalisation-dualité d'une classe de problèmes de dirichlet non linéaires. *ESAIM: Mathematical Modelling and Numerical Analysis*, 9(R2):41–76, 1975. 3

[17] D. Han and X. Yuan. A note on the alternating direction method of multipliers. *Journal of Optimization Theory and Applications*, 155(1):227–238, 2012. 3

[18] B. He, L. Hou, and X. Yuan. On full Jacobian decomposition of the augmented Lagrangian method for separable convex programming. *SIAM Journal on Optimization*, 25(4):2274–2312, 2015. 4

[19] B. He, M. Tao, and X. Yuan. Alternating direction method with gaussian back substitution for separable convex programming. *SIAM Journal on Optimization*, 22(2):313–340, 2012. 3, 4, 16

[20] B. He, M. Tao, and X. Yuan. Convergence rate and iteration complexity on the alternating direction method of multipliers with a substitution procedure for separable convex programming. *Optimization Online*, 2012. 4

[21] B. He and X. Yuan. On the $O(1/n)$ convergence rate of the Douglas-Rachford alternating direction method. *SIAM Journal on Numerical Analysis*, 50(2):700–709, 2012. 3

[22] C. Hildreth. A quadratic programming procedure. *Naval Research Logistics Quarterly*, 4(1):79–85, 1957. 3

[23] M. Hong, T.-H. Chang, X. Wang, M. Razaviyayn, S. Ma, and Z.-Q. Luo. A block successive upper bound minimization method of multipliers for linearly constrained convex optimization. *arXiv preprint arXiv:1401.7079*, 2014. 3, 4

[24] M. Hong, X. Wang, M. Razaviyayn, and Z.-Q. Luo. Iteration complexity analysis of block coordinate descent methods. *arXiv preprint arXiv:1310.6957v2*, 2015. 3

[25] R. A. Horn and C. R. Johnson. Topics in matrix analysis. *Cambridge UP, New York*, 1991. 7

[26] G. M. James, C. Paulson, and P. Rusmevichientong. The constrained lasso. Technical report, Citeseer, 2012. 2

[27] M. Li, D. Sun, and K.-C. Toh. A convergent 3-block semi-proximal ADMM for convex minimization problems with one strongly convex block. *Asia-Pacific Journal of Operational Research*, 32(04):1550024, 2015. 3

[28] X. Li, D. Sun, and K.-C. Toh. A schur complement based semi-proximal admm for convex quadratic conic programming and extensions. *Mathematical Programming*, 155(1-2):333–373, 2016. 3, 4

[29] T. Lin, S. Ma, and S. Zhang. On the global linear convergence of the admm with multiblock variables. *SIAM Journal on Optimization*, 25(3):1478–1497, 2015. 3

[30] T. Lin, S. Ma, and S. Zhang. On the sublinear convergence rate of multi-block admm. *Journal of the Operations Research Society of China*, 3(3):251–274, 2015. 3

[31] R. D. Monteiro and B. F. Svaiter. Iteration-complexity of block-decomposition algorithms and the alternating direction method of multipliers. *SIAM Journal on Optimization*, 23(1):475–507, 2013. 3

[32] I. Necoara and A. Patrascu. A random coordinate descent algorithm for optimization problems with composite objective function and linear coupled constraints. *Computational Optimization and Applications*, 57(2):307–337, 2014. 3

[33] Z. Peng, T. Wu, Y. Xu, M. Yan, and W. Yin. Coordinate friendly structures, algorithms and applications. *Annals of Mathematical Sciences and Applications*, 1(1):57–119, 2016. 2, 3

[34] Z. Peng, Y. Xu, M. Yan, and W. Yin. Arock: an algorithmic framework for asynchronous parallel coordinate updates. *arXiv preprint arXiv:1506.02396*, 2015. 3

[35] J.-C. Pesquet and A. Repetti. A class of randomized primal-dual algorithms for distributed optimization. *arXiv preprint arXiv:1406.6404*, 2014. 3

[36] M. Razaviyayn, M. Hong, and Z.-Q. Luo. A unified convergence analysis of block successive minimization methods for nonsmooth optimization. *SIAM Journal on Optimization*, 23(2):1126–1153, 2013. 3

[37] D. Sun, K.-C. Toh, and L. Yang. A convergent 3-block semiproximal alternating direction method of multipliers for conic programming with 4-type constraints. *SIAM journal on Optimization*, 25(2):882–915, 2015. 3, 4

[38] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996. 2

[39] K.-C. Toh, M. J. Todd, and R. H. Tütüncü. SDPT3—a matlab software package for semidefinite programming, version 1.3. *Optimization methods and software*, 11(1-4):545–581, 1999. 15

[40] P. Tseng. Convergence of a block coordinate descent method for nondifferentiable minimization. *Journal of Optimization Theory and Applications*, 109(3):475–494, 2001. 3

[41] P. Tseng and S. Yun. Block-coordinate gradient descent method for linearly constrained nonsmooth separable optimization. *Journal of optimization theory and applications*, 140(3):513–535, 2009. 3

[42] P. Tseng and S. Yun. A coordinate gradient descent method for nonsmooth separable minimization. *Mathematical Programming*, 117(1-2):387–423, 2009. 3

[43] Y. Wang, W. Yin, and J. Zeng. Global convergence of ADMM in nonconvex nonsmooth optimization. *arXiv preprint arXiv:1511.06324*, 2015. 3

[44] J. Warga. Minimizing certain convex functions. *Journal of the Society for Industrial and Applied Mathematics*, 11(3):588–593, 1963. 3

[45] J. Wright, A. Ganesh, K. Min, and Y. Ma. Compressive principal component pursuit. *Information and Inference*, 2(1):32–68, 2013. 3, 18

[46] S. J. Wright. Coordinate descent algorithms. *Mathematical Programming*, 151(1):3–34, 2015. 3

[47] Y. Xu. Accelerated first-order primal-dual proximal methods for linearly constrained composite convex programming. *arXiv preprint arXiv:1606.09155*, 2016. 7

[48] Y. Xu, I. Akrotirianakis, and A. Chakraborty. Alternating direction method of multipliers for regularized multiclass support vector machines. In *International Workshop on Machine Learning, Optimization and Big Data*, pages 105–117. Springer, 2015. 3, 19

[49] Y. Xu, I. Akrotirianakis, and A. Chakraborty. Proximal gradient method for huberized support vector machine. *Pattern Analysis and Applications*, pages 1–17, 2015. 20

[50] Y. Xu and W. Yin. A block coordinate descent method for regularized multiconvex optimization with applications to nonnegative tensor factorization and completion. *SIAM Journal on imaging sciences*, 6(3):1758–1789, 2013. 3

[51] Y. Xu and W. Yin. A globally convergent algorithm for nonconvex optimization based on block coordinate update. *arXiv preprint arXiv:1410.1386*, 2014. 3